L4 – Domain Name System (DNS)

by
T.S.R.K. Prasad

EA C451 Internetworking Technologies

19/01/2013

References / Acknowledgements

Sec 2.5: DNS, [Kurose]

Sec 9.3.1: Name Service (DNS), [Peterson]

Sec 12.3: DNS, [Farrel]

Optional Readings

[rfc920] - Domain Requirements

[rfc1101] - DNS Encoding of Network Names and Other Types

[rfc1034] - Domain Names – Concepts and Facilities

[rfc1035] - Domain Names – Implementation and Specification

[rfc1546] - Host Anycasting Service

[rfc4367] - What's in a Name: False Assumptions about DNS Names

Optional Readings

[rfc4592] - Wildcards in dns

[rfc5782] - DNS Blacklists and Whitelists

[Weaver] Nicholas Weaver, Christian Kreibich, and Vern Paxson, **Redirecting DNS for Ads and Profit**

[Walfish-SFR] Michael Walfish, Hari Balakrishnan, and Scott Shenker, **Untangling the Web from DNS**

Presentation Overview

- Problem Areas
- Queries and Responses
- Resource Records
- Name Space
- DNS Basics

Presentation Overview

Problem Areas

Queries and Responses

Resource Records

Name Space

DNS Basics

Host Names vs. IP addresses

- **Host names**
  - Mnemonic name appreciated by humans
  - Variable length, full alphabet of characters
  - Provide little (if any) information about location
  - Examples: www.cnn.com and bbc.co.uk
- **IP addresses**
  - Numerical address appreciated by routers
  - Fixed length, binary number
  - Hierarchical, related to host location
  - Examples: 64.236.16.20 and 212.58.228.155

Separating Names and Addresses

- **Names are easier to remember**
  - www.cnn.com vs. 64.236.16.20
- **Addresses can change underneath**
  - Move www.cnn.com to 64.236.16.20
  - E.g., renumbering when changing providers
- **Name could map to multiple IP addresses**
  - www.cnn.com to multiple (8) replicas of the Web site
- **Map to different addresses* in different places**
  - Address of a nearby copy of the Web site
  - E.g., to reduce latency, or return different content
- **Multiple names for the same address**
- E.g., aliases like www.cnn.com and cnn.com

Scalable (Name ↔ Address) Mappings

- **Originally: per-host file**
  - Flat namespace
  - **/etc/hosts** (what is this on your computer today?)
  - SRI kept master copy
  - Downloaded regularly
- **Single server doesn't scale**
  - Traffic implosion (lookups & updates)
  - Single point of failure
  - Amazing politics

Need a distributed, hierarchical collection of servers

Properties of Present DNS

- **Hierarchical namespace**
  - Hierarchical name space divided into zones
- **Hierarchy of DNS servers**
  - Root (hardwired into other servers)
  - Top-level domain (TLD) servers
  - Authoritative DNS servers
  - Zones distributed over collection of DNS servers
- **Performing the translations**
  - Local DNS servers
  - Resolver software

Presentation Overview

Problem Areas

Queries and Responses

Resource Records

Name Space

DNS Basics

DNS Name Space

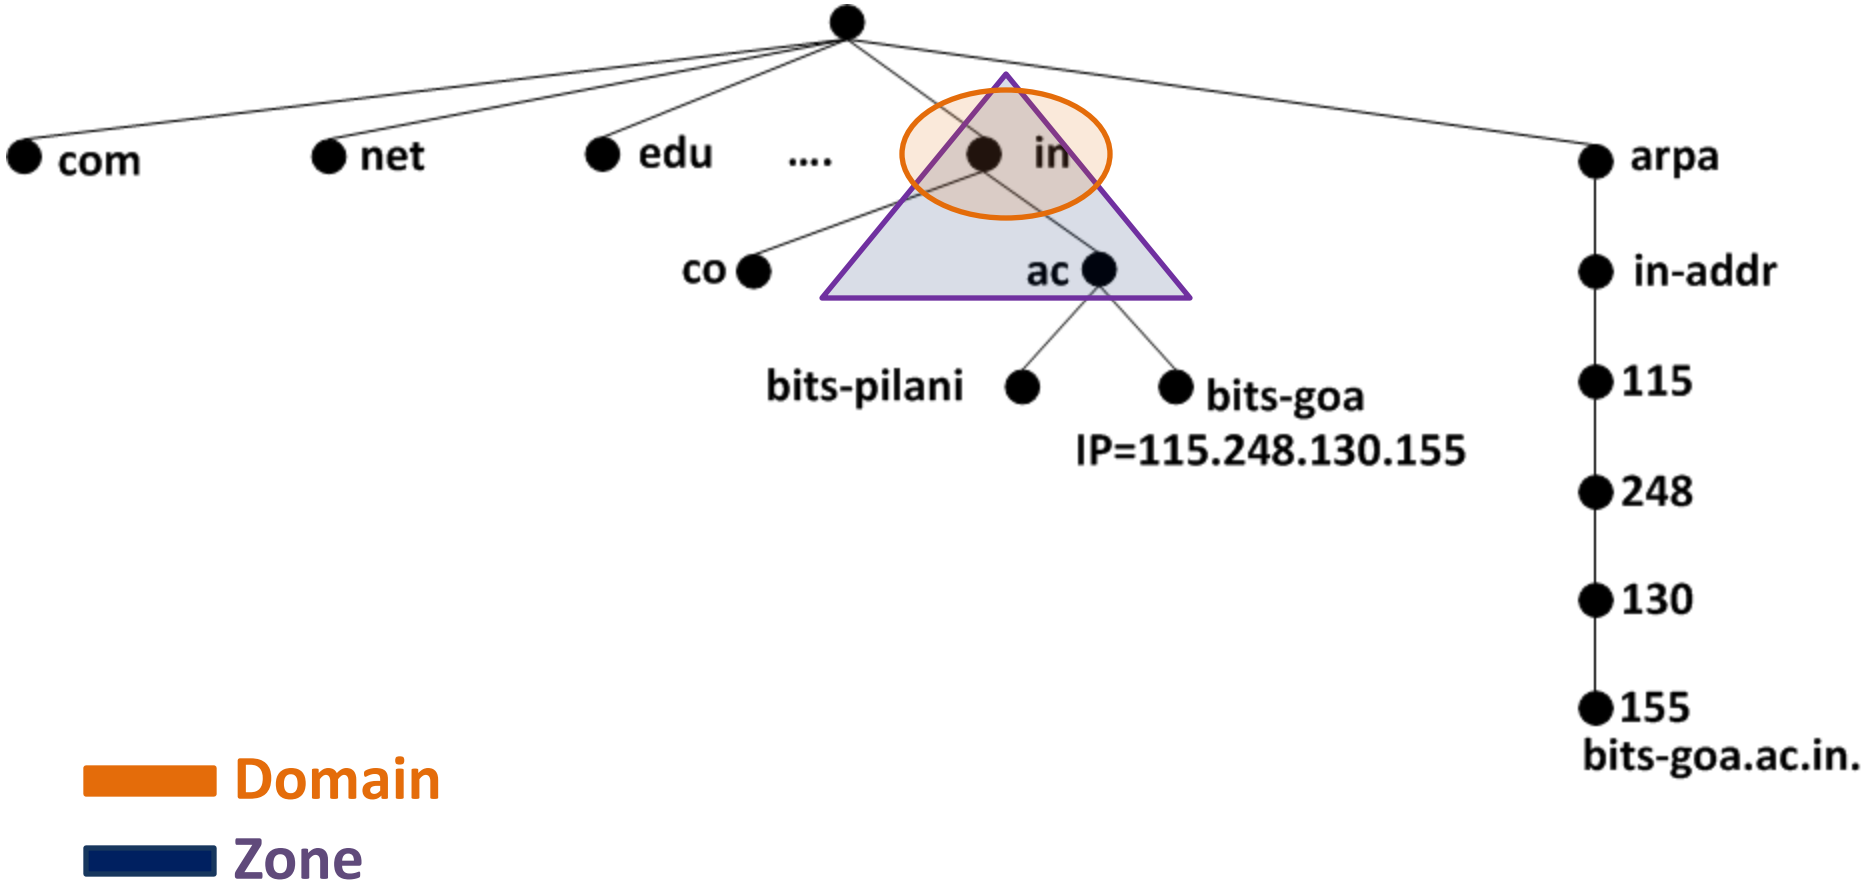

Root Servers

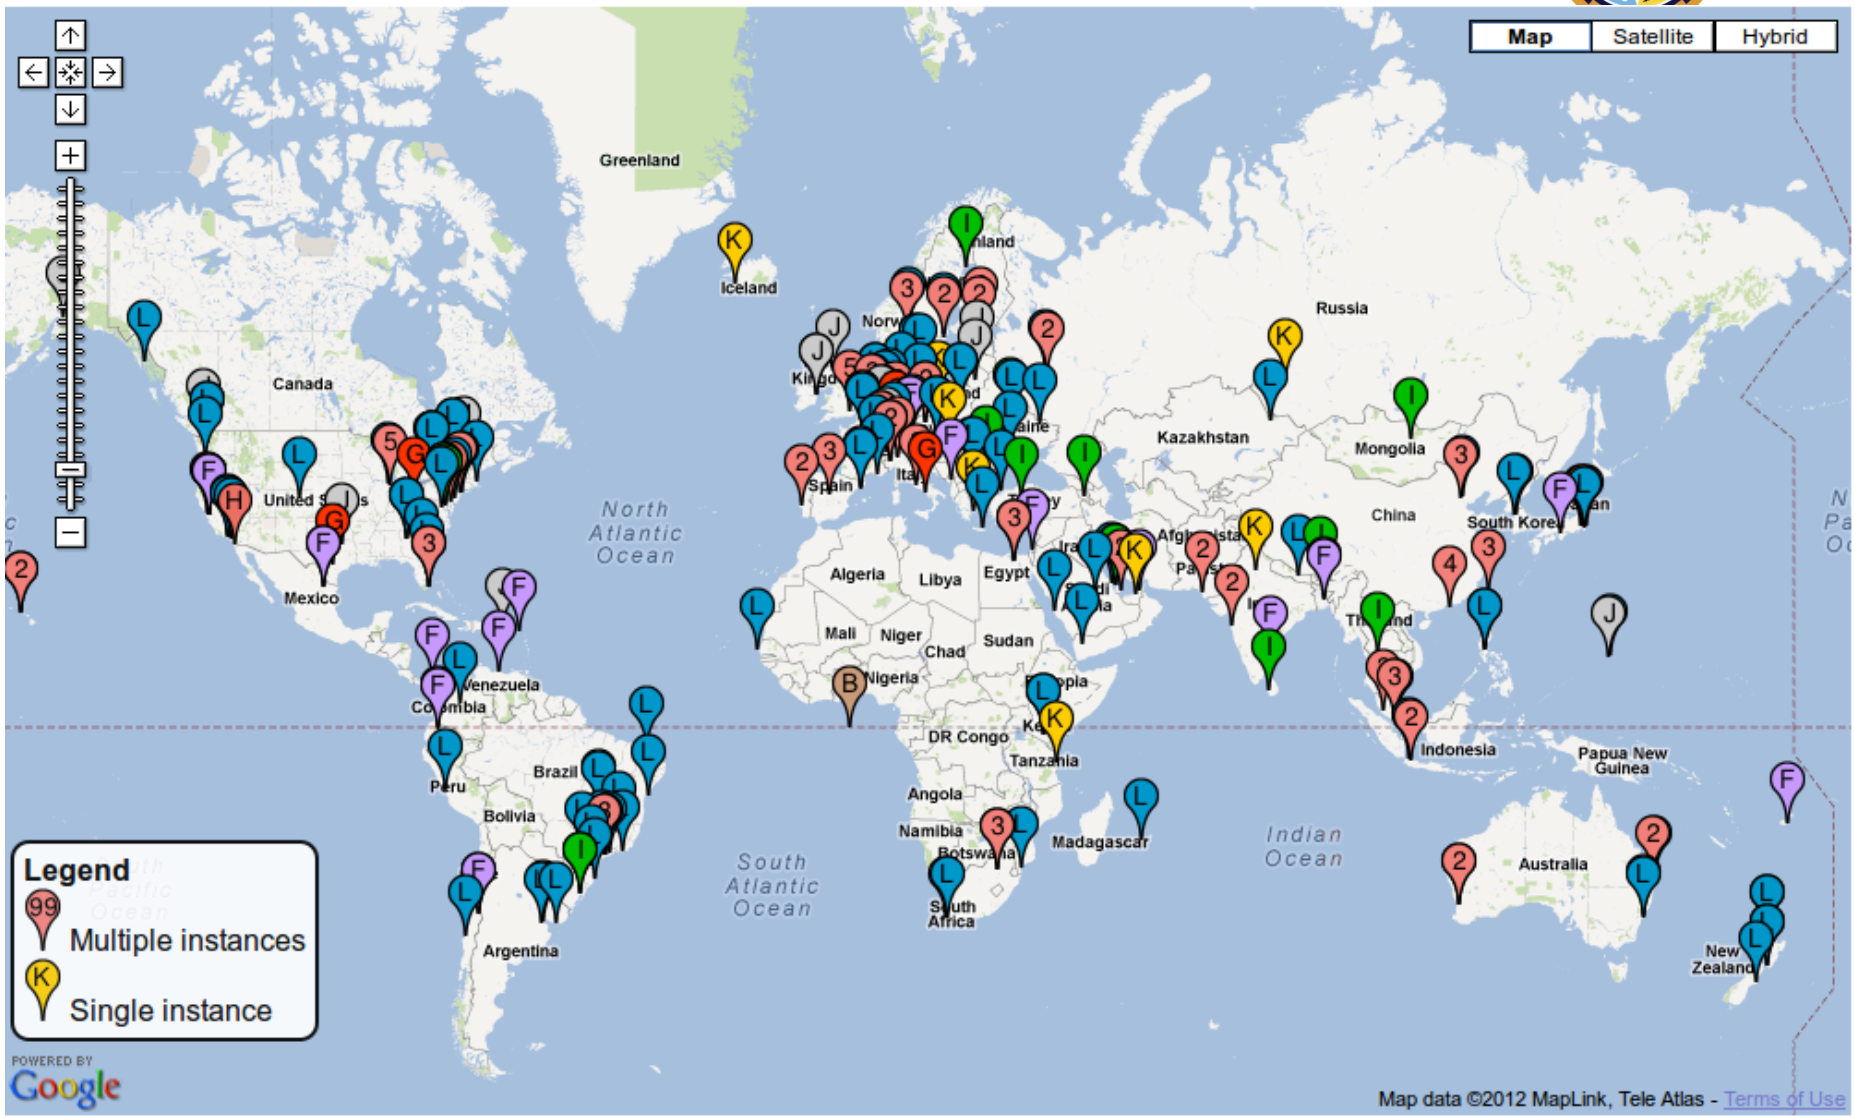

Root Servers in India

gTLD	Description
aero	Airlines and aerospace companies
biz	Businesses or firms
com	Commercial organizations
coop	Cooperative business organizations
edu	Educational institutions
gov	Government institutions
info	Information service providers
int	International organizations
mil	Military groups
museum	Museums
name	Personal names / Individuals
net	Network support centers
org	Nonprofit organizations
pro	Profesional individual organizations

Generic Top Level Domains (gTLDs)

TLD and Authoritative DNS Servers

- **Top-level domain (TLD) servers**
  - Generic domains (e.g., com, org, edu)
  - Country domains (e.g., uk, fr, cn, jp)
  - Special domains (e.g., arpa)
  - Typically managed professionally
  - Network Solutions maintains servers for "**com**"
  - Educause maintains servers for "**edu**"
- **Authoritative DNS servers**
  - Provide public records for hosts at an organization
  - Private records may differ, though not part of original design's intent
  - For the organization's servers (e.g., Web and mail)
  - Can be maintained locally or by a service provider

Local DNS Server

- does not strictly belong to hierarchy
- each ISP (residential ISP, company, university) has one
  - also called “default name server”
- when host makes DNS query, query is sent to its local DNS server
  - has local cache of recent name-to-address translation pairs (but may be out of date!)
  - acts as proxy, forwards query into hierarchy

Presentation Overview

Problem Areas

Queries and Responses

Resource Records

Name Space

DNS Basics

Resource Records (RR)

Format

<Name, Type, Value(Rlength+Rdata), Class, TTL>

Common Types:

A

NS

MX

PTR

(many others) → ***A6, AAAA, SOA, CNAME, TXT, SPF, SRV, HINFO, DNAME, GPOS, LOC, NSAP, KEY, …***

Root Server Zone File (in 11/1987)


```
.       IN      SOA     SRI-NIC.ARPA. HOSTMASTER.SRI-NIC.ARPA. (
                        870611          ;serial
                        1800            ;refresh every 30 min
                        300             ;retry every 5 min
                        604800          ;expire after a week
                        86400)          ;minimum of a day
                NS      A.ISI.EDU.
                NS      C.ISI.EDU.
                NS      SRI-NIC.ARPA.

MIL.    86400   NS      SRI-NIC.ARPA.
        86400   NS      A.ISI.EDU.
```

Root Server Zone File (in 11/1987)


```
EDU.    86400   NS      SRI-NIC.ARPA.
        86400   NS      C.ISI.EDU.

SRI-NIC.ARPA.   A       26.0.0.73
                A       10.0.0.51
                MX      0 SRI-NIC.ARPA.
                HINFO   DEC-2060 TOPS20

ACC.ARPA.       A       26.6.0.65
                HINFO   PDP-11/70 UNIX
                MX      10 ACC.ARPA.

USC-ISIC.ARPA.  CNAME   C.ISI.EDU.
```

```
73.0.0.26.IN-ADDR.ARPA.  PTR   SRI-NIC.ARPA.
65.0.6.26.IN-ADDR.ARPA.  PTR   ACC.ARPA.
51.0.0.10.IN-ADDR.ARPA.  PTR   SRI-NIC.ARPA.
52.0.0.10.IN-ADDR.ARPA.  PTR   C.ISI.EDU.

103.0.3.26.IN-ADDR.ARPA. PTR   A.ISI.EDU.

A.ISI.EDU. 86400 A     26.3.0.103
C.ISI.EDU. 86400 A     10.0.0.52
```

From ARPA. Zone File …

arpa.	518400	IN	NS	a.root-servers.net.
…				
arpa.	518400	IN	NS	m.root-servers.net.
in-addr.arpa.	172800	IN	NS	a.in-addr-servers.arpa.
…				
in-addr.arpa.	172800	IN	NS	f.in-addr-servers.arpa.
a.in-addr-servers.arpa.	172800	IN	A	199.212.0.73
a.in-addr-servers.arpa.	172800	IN	AAAA	2001:500:13:0:0:0:0:73
…				
f.in-addr-servers.arpa.	172800	IN	A	193.0.9.1
f.in-addr-servers.arpa.	172800	IN	AAAA	2001:67c:e0:0:0:0:0:1

ARPA - Address and Routing Parameter Area Domain

Presentation Overview

Problem Areas

Queries and Responses

Resource Records

Name Space

DNS Basics

DNS Message Format

<DNS Message>::= <common header>

<question> [<question> . . .]

[<answer>. . .]

[<authority> . . .]

[<additional info> . . .]

DNS Header Format

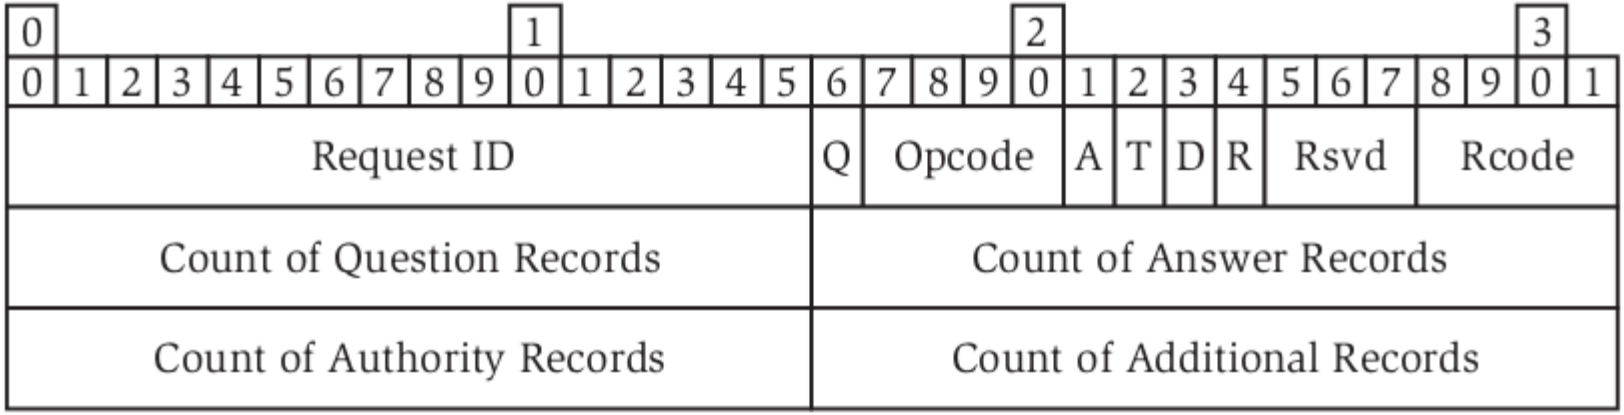

Opcode	*Request Type*
0	A standard query
1	An inverse query
2	A server status request

RCode	Meaning
0	Success / Message Request
1	Format Error
2	Server Failure
3	Name Error
4	Not Implemented
5	Refused

DNS Header Message Bits

Bit	Meaning
Q	Query
A	Authoritative Answer
T	Truncated Response
D	Recursion Desired
R	Recursion Available

DNS Record Formats

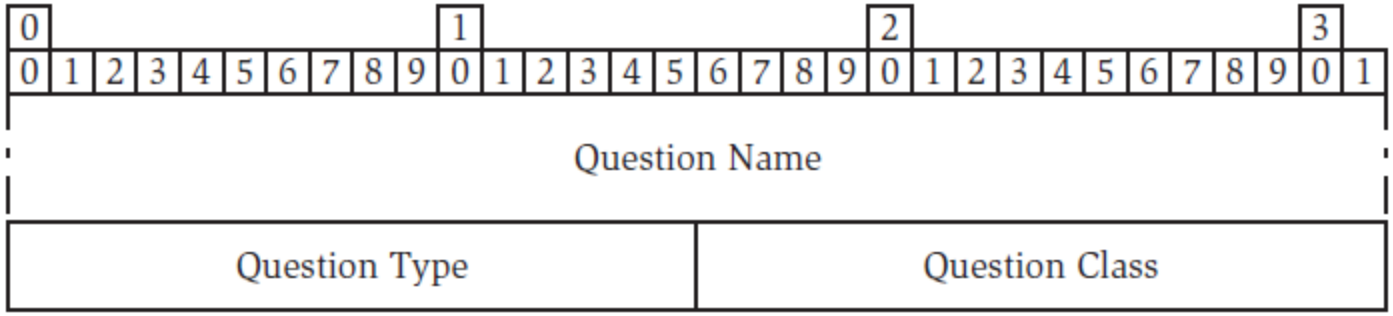

Question Record

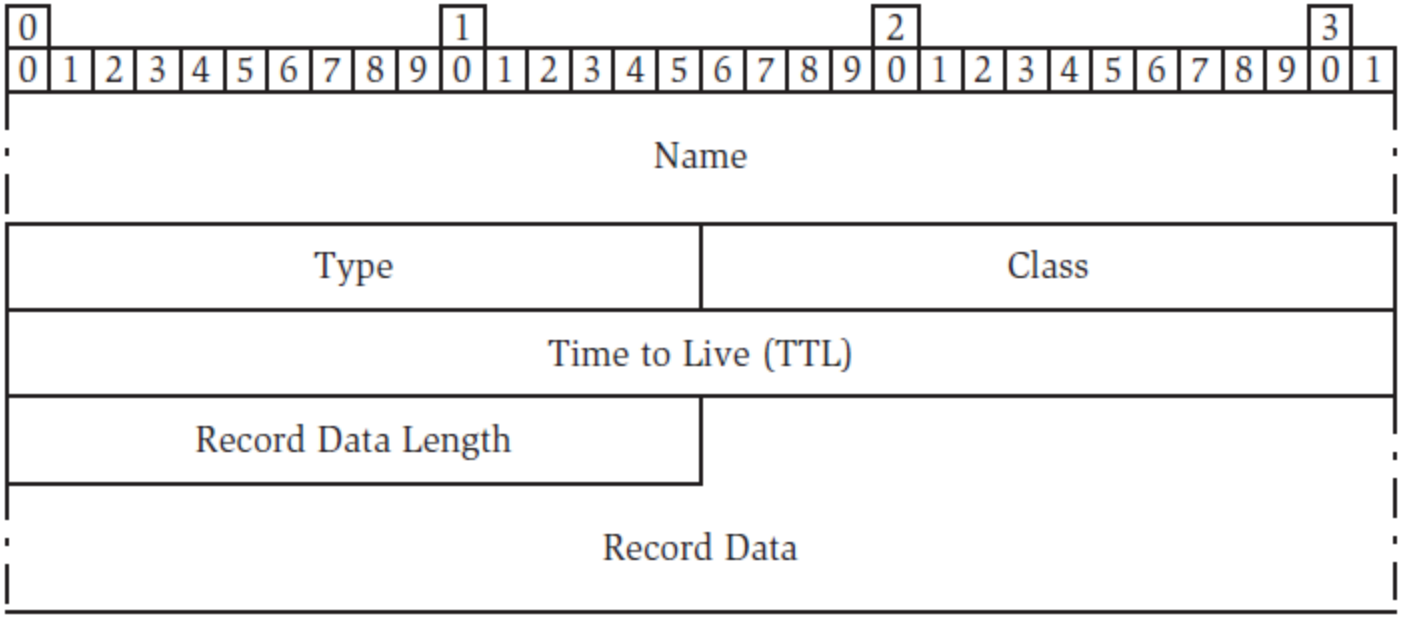

Response Record

Iterative Query

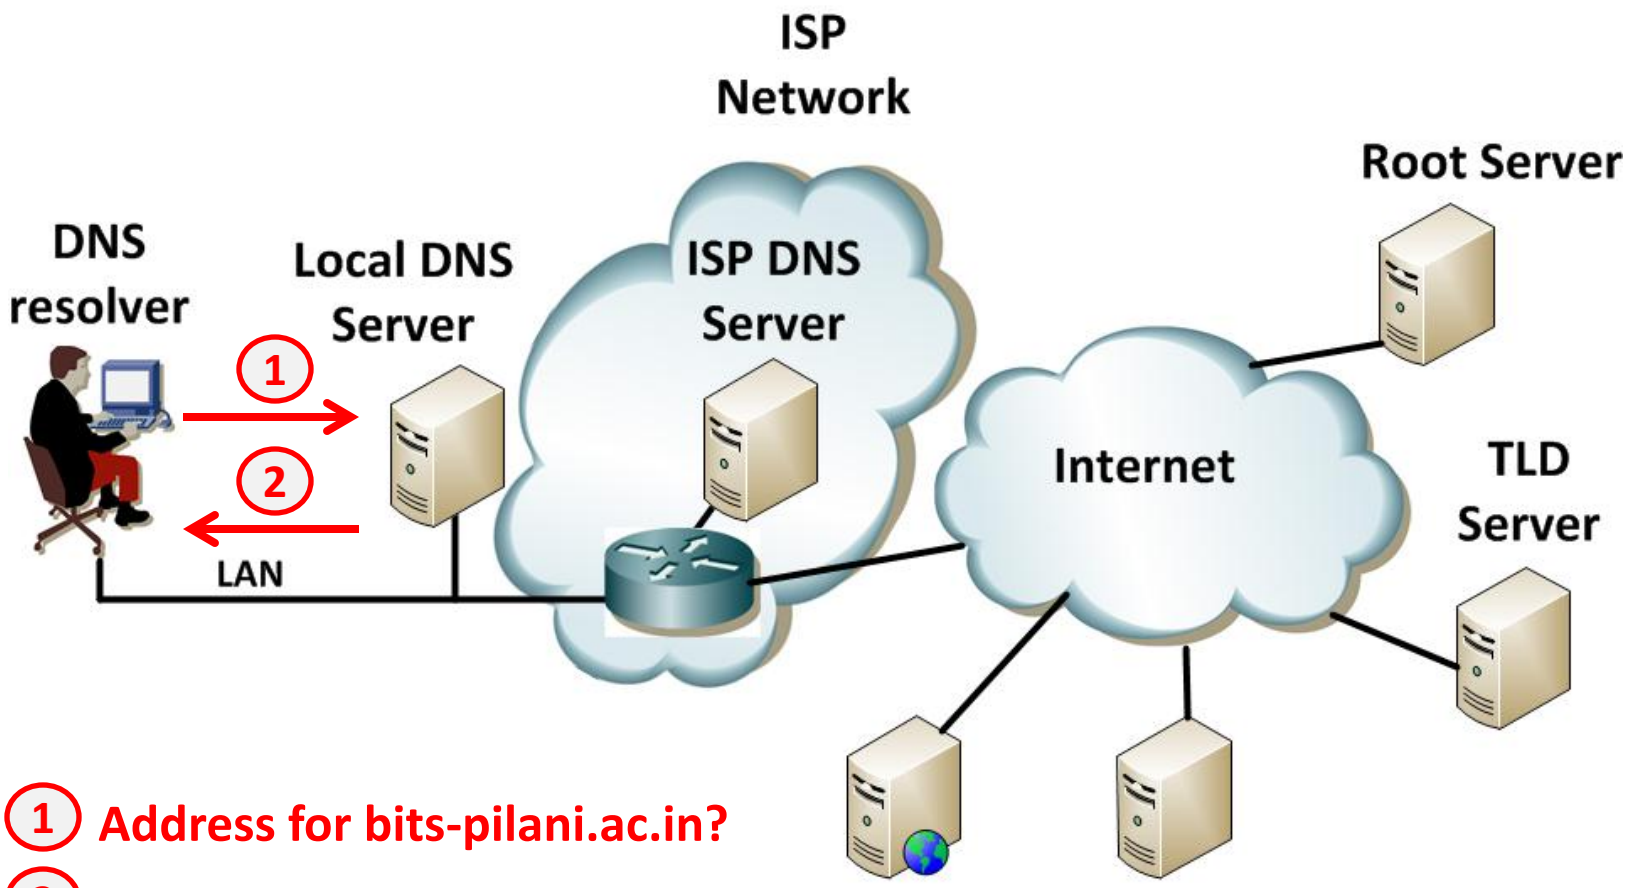

Serves both ***.in.*** and ***.ac.in.*** domains

1. **Address for bits-pilani.ac.in?**
2. **Don't know; Ask ISP DNS Server**

Iterative Query

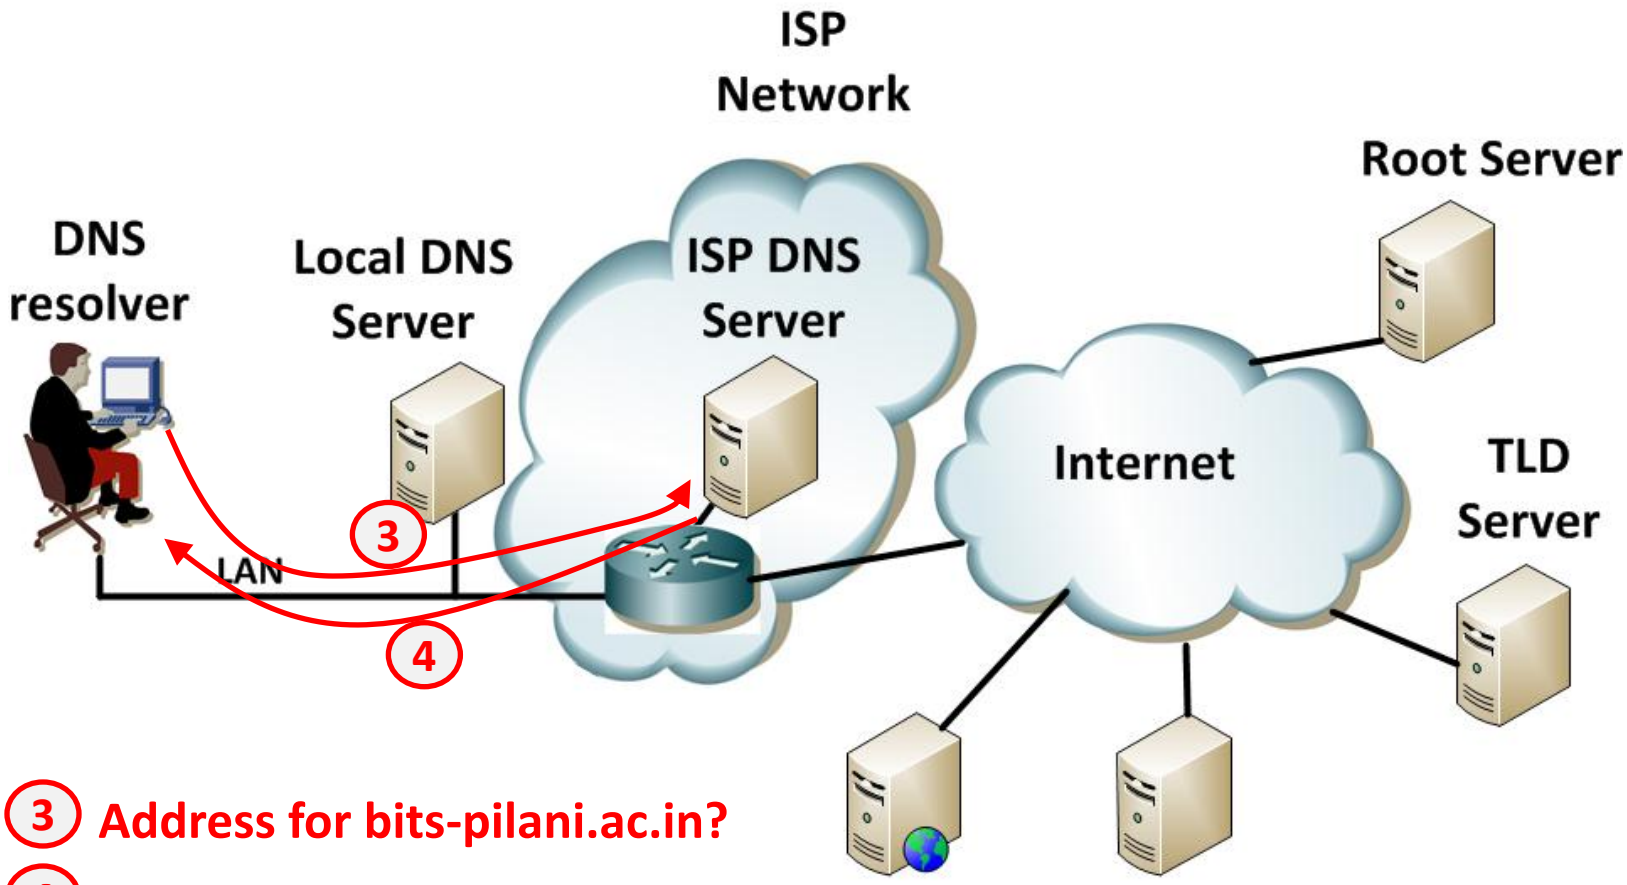

3 **Address for bits-pilani.ac.in?**

4 **Don't know; Ask the Root Server**

Serves both ***.in.*** and ***.ac.in.*** domains

Iterative Query

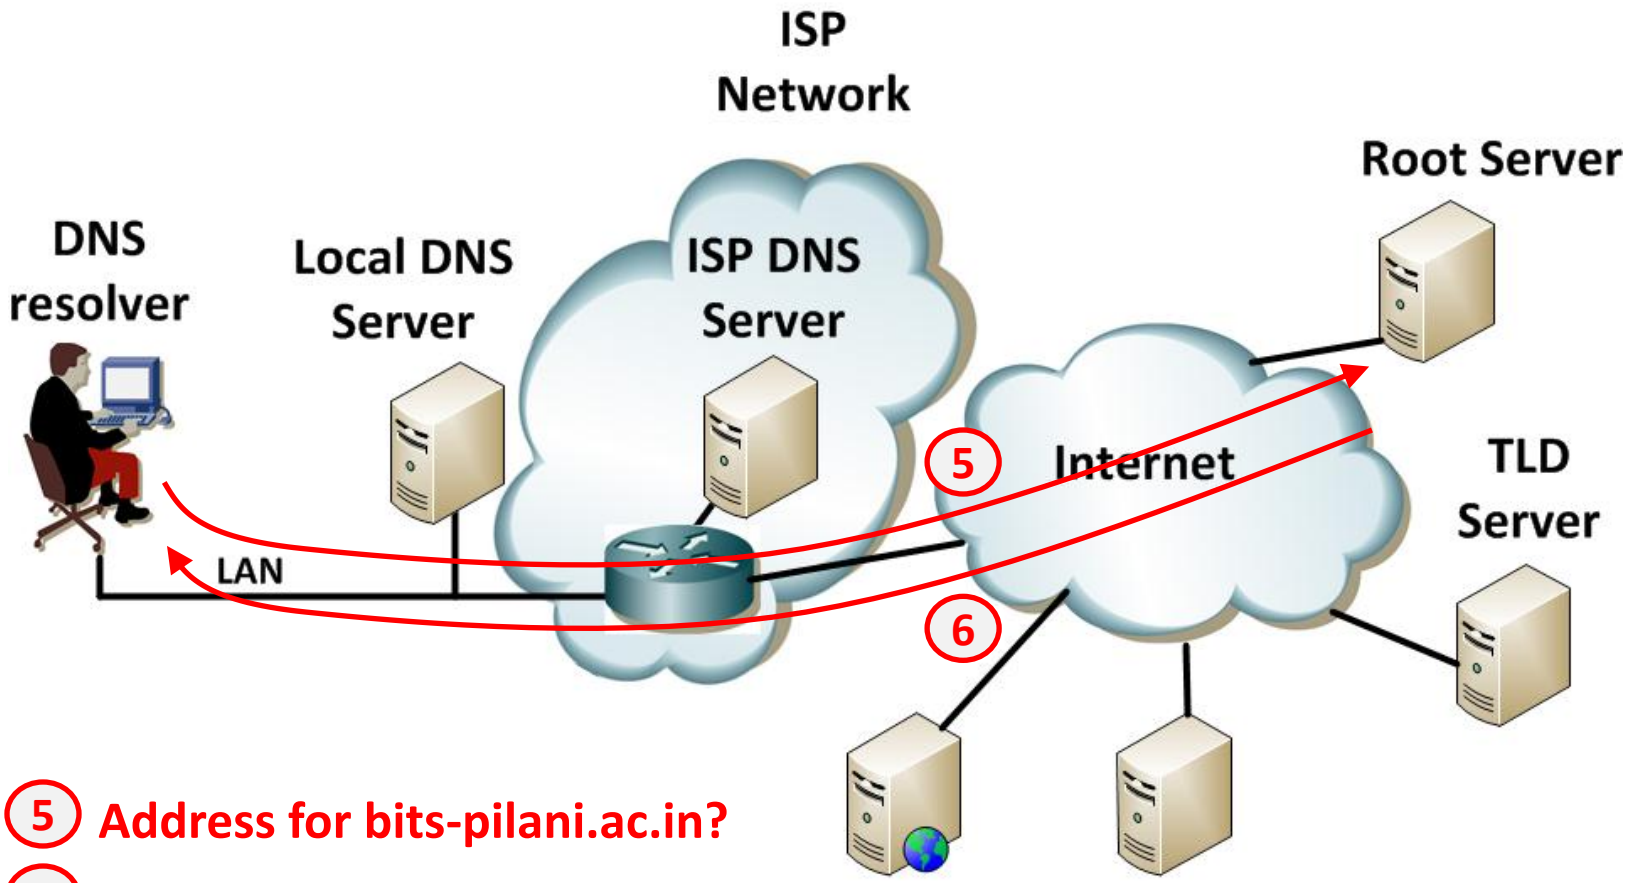

5 Address for bits-pilani.ac.in?

6 Don't know; Ask TLD Server

Serves both ***.in.*** and ***.ac.in.*** domains

Iterative Query

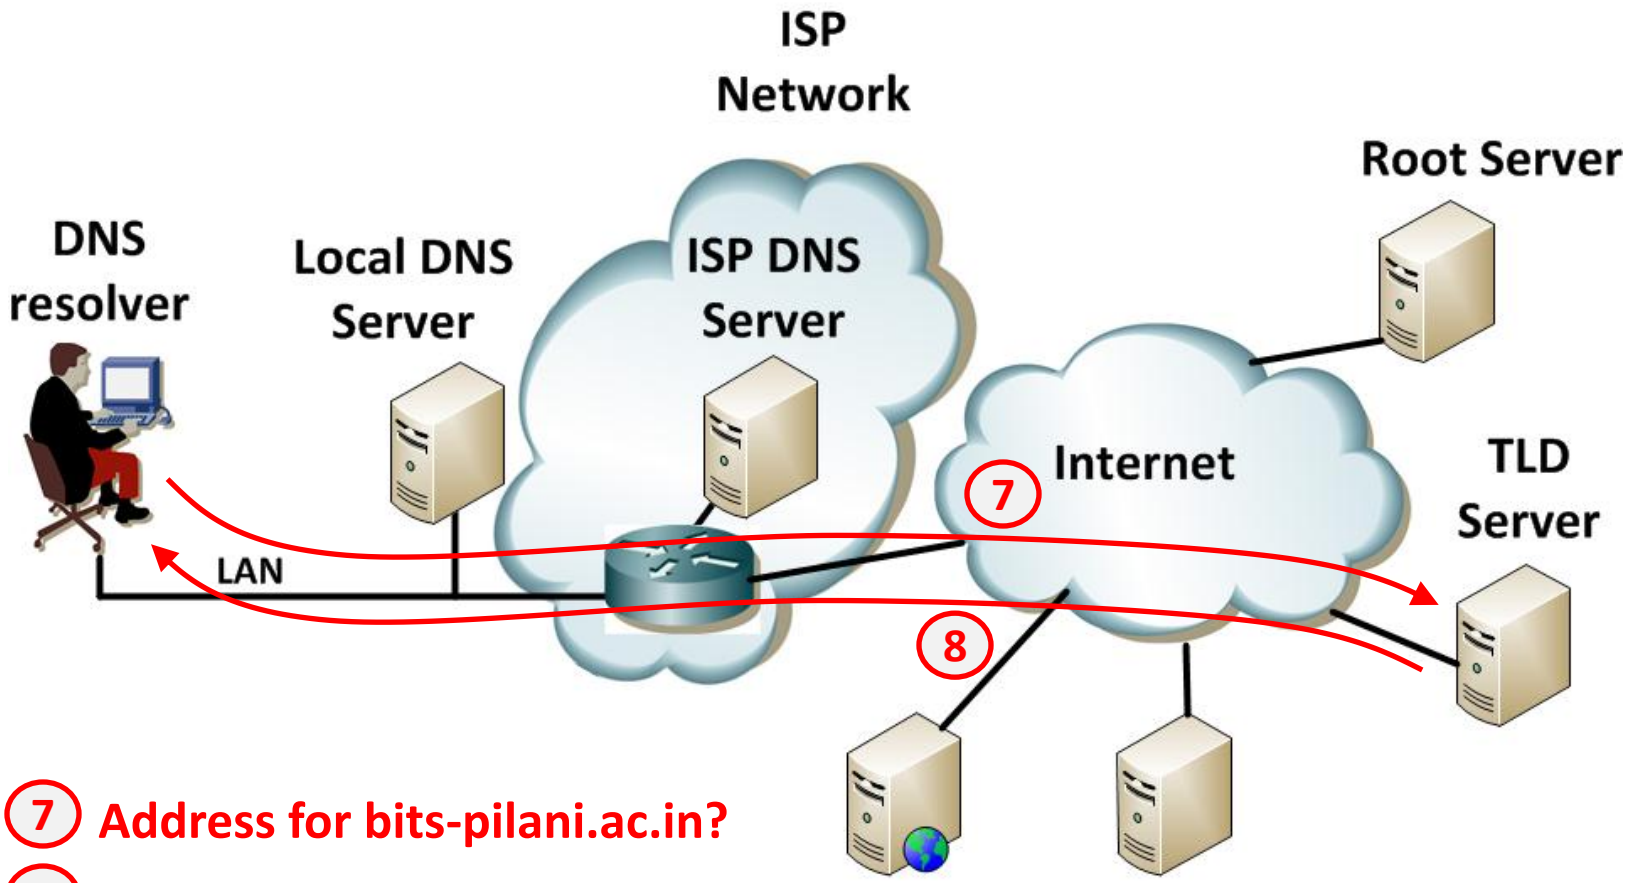

Serves both ***.in.*** and ***.ac.in.*** domains

7 **Address for bits-pilani.ac.in?**

8 **Don't know; Ask the authoritative name server**

Iterative Query

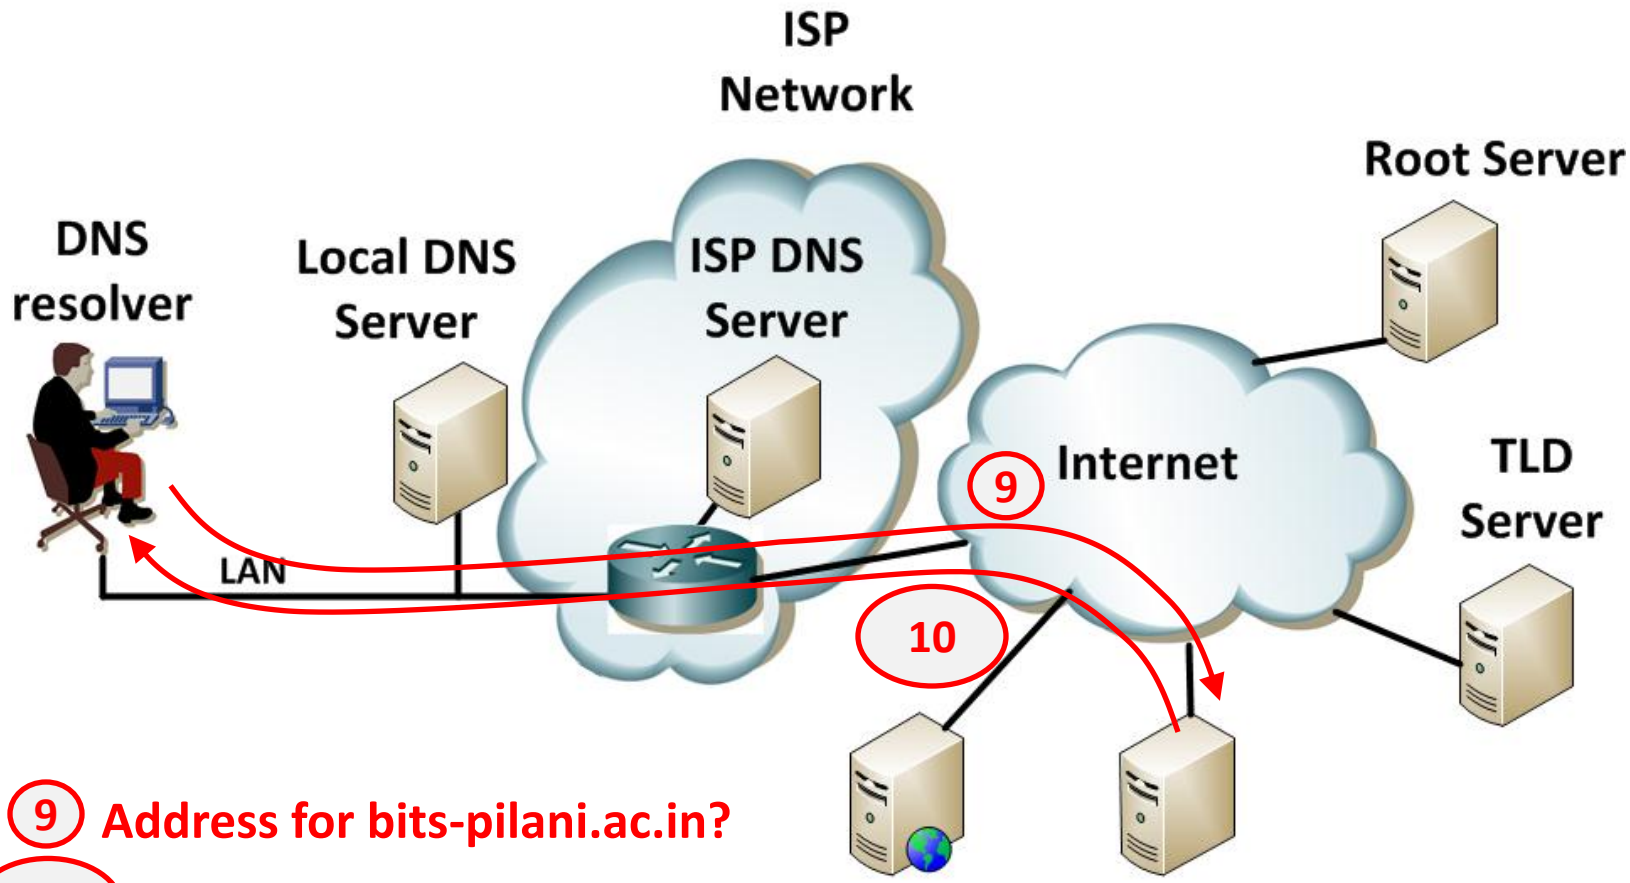

Serves both ***.in.*** and ***.ac.in.*** domains

9 **Address for bits-pilani.ac.in?**

10 **Address is 10.10.10.100**

Recursive Query

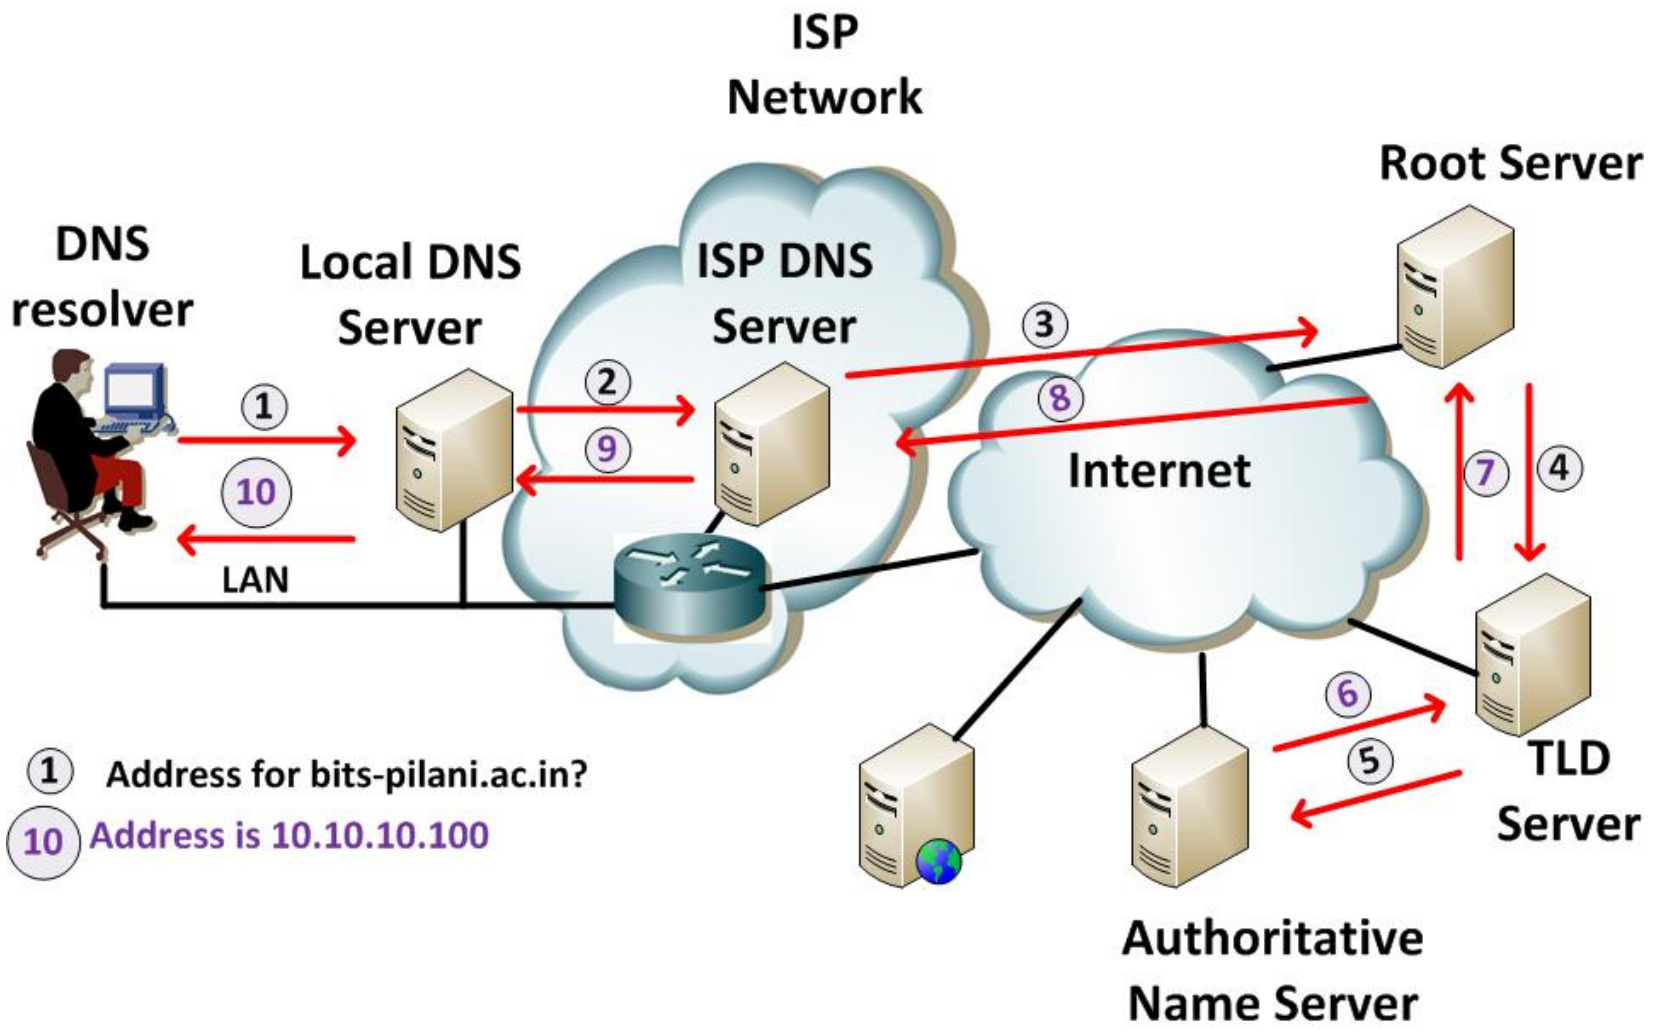

Serves both ***.in.*** and ***.ac.in.*** domains

Types of DNS Queries

- **Forward DNS Query**
  - Domain name to other resources
  - Uses dedicated forward DNS infrastructure
- **Reverse DNS Query**
  - IP address to domain name
  - Uses dedicated reverse DNS infrastructure
- **Inverse DNS Query**
  - Resources to domain name
  - Need to know the authoritative name server

Sample Query

Header	OPCODE=SQUERY
Question	QNAME=SRI-NIC.ARPA., QCLASS=IN, QTYPE=A
Answer	<empty>
Authority	<empty>
Additional	<empty>

Response to Sample Query

Header	OPCODE=SQUERY, RESPONSE, AA
Question	QNAME=SRI-NIC.ARPA., QCLASS=IN, QTYPE=A
Answer	SRI-NIC.ARPA. 86400 IN A 26.0.0.73 86400 IN A 10.0.0.51
Authority	<empty>
Additional	<empty>

Response for Another Query

Answer for question: **QNAME=BRL.MIL, QTYPE=A**

Header	OPCODE=SQUERY, RESPONSE				
Question	QNAME=BRL.MIL, QCLASS=IN, QTYPE=A				
Answer	<empty>				
Authority	MIL.	86400	IN	NS	SRI-NIC.ARPA.
		86400		NS	A.ISI.EDU.
Additional	A.ISI.EDU.			A	26.3.0.103
	SRI-NIC.ARPA.			A	26.0.0.73
				A	10.0.0.51

Reverse DNS Query for 68.142.214.24 (flickr.com)

The ordered list of name servers used for the reverse query -

1. NS: b.root-servers.net [192.228.79.201] (Marina Del Rey, California, United States)
2. NS: c.in-addr-servers.arpa [196.216.169.10] (Mauritius)
3. NS: y.arin.net [192.42.93.32] (Sterling, Virginia, United States)
4. NS: ns3.yahoo.com [121.101.152.99] (Mumbai, Maharashtra, India)

Answer:

24.214.142.68.in-addr.arpa. 1800 undefined PTR www.flickr.vip.mud.yahoo.com

DNS Caching

- **Performing all these queries takes time**
  - And all this before actual communication takes place
  - E.g., 1-second latency before starting Web download
- **Caching can greatly reduce overhead**
  - The top-level servers very rarely change
  - Popular sites (e.g., www.cnn.com) visited often
  - Local DNS server often has the information cached
- **How DNS caching works**
  - DNS servers cache responses to queries
  - Responses include a “time to live” (TTL) field
  - Server deletes cached entry after TTL expires

Presentation Overview

Problem Areas

Queries and Responses

Resource Records

Name Space

DNS Basics

DNSRedirect Using Wildcards

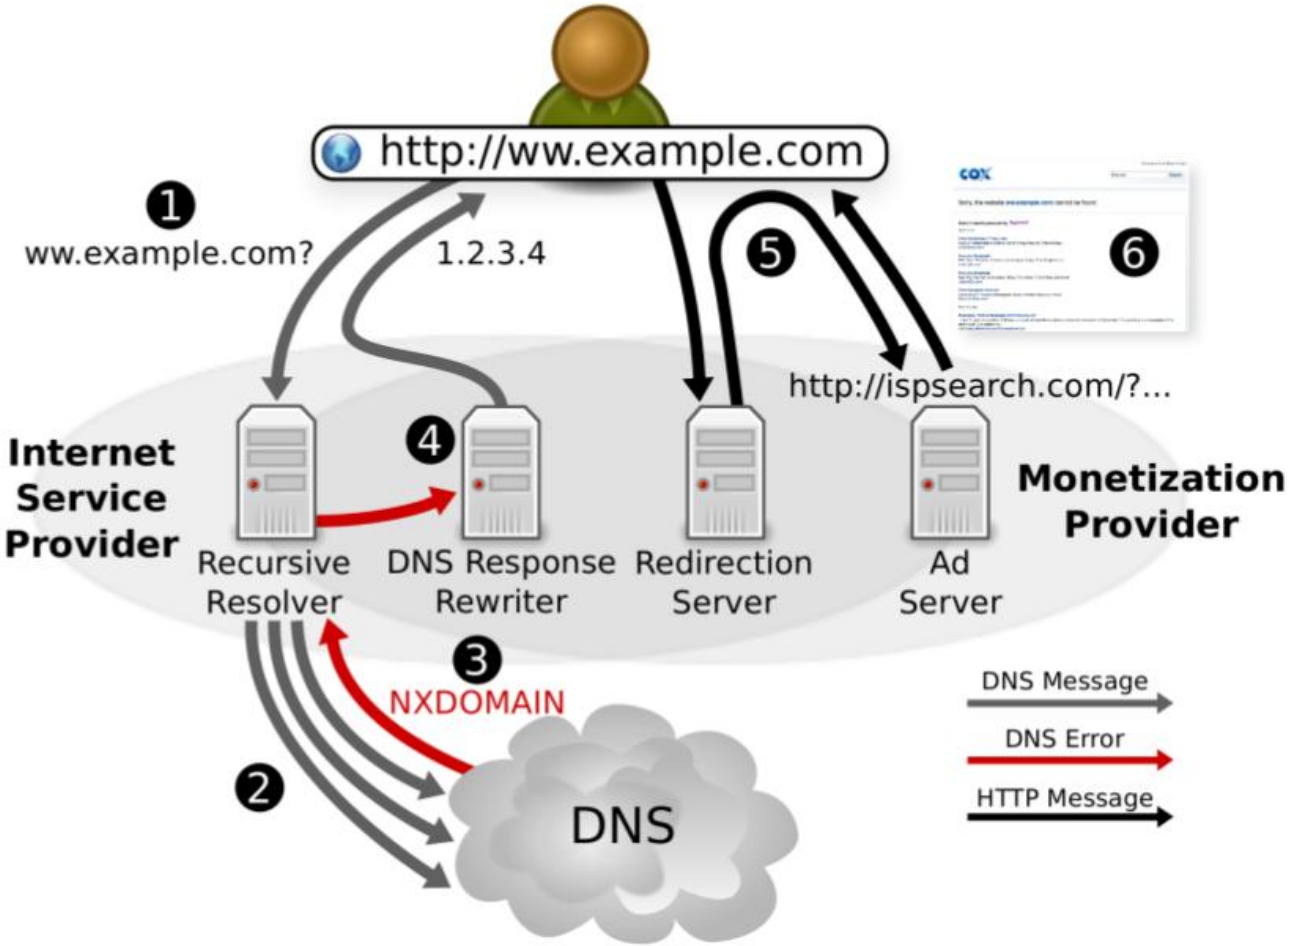

Result of DNS Redirect

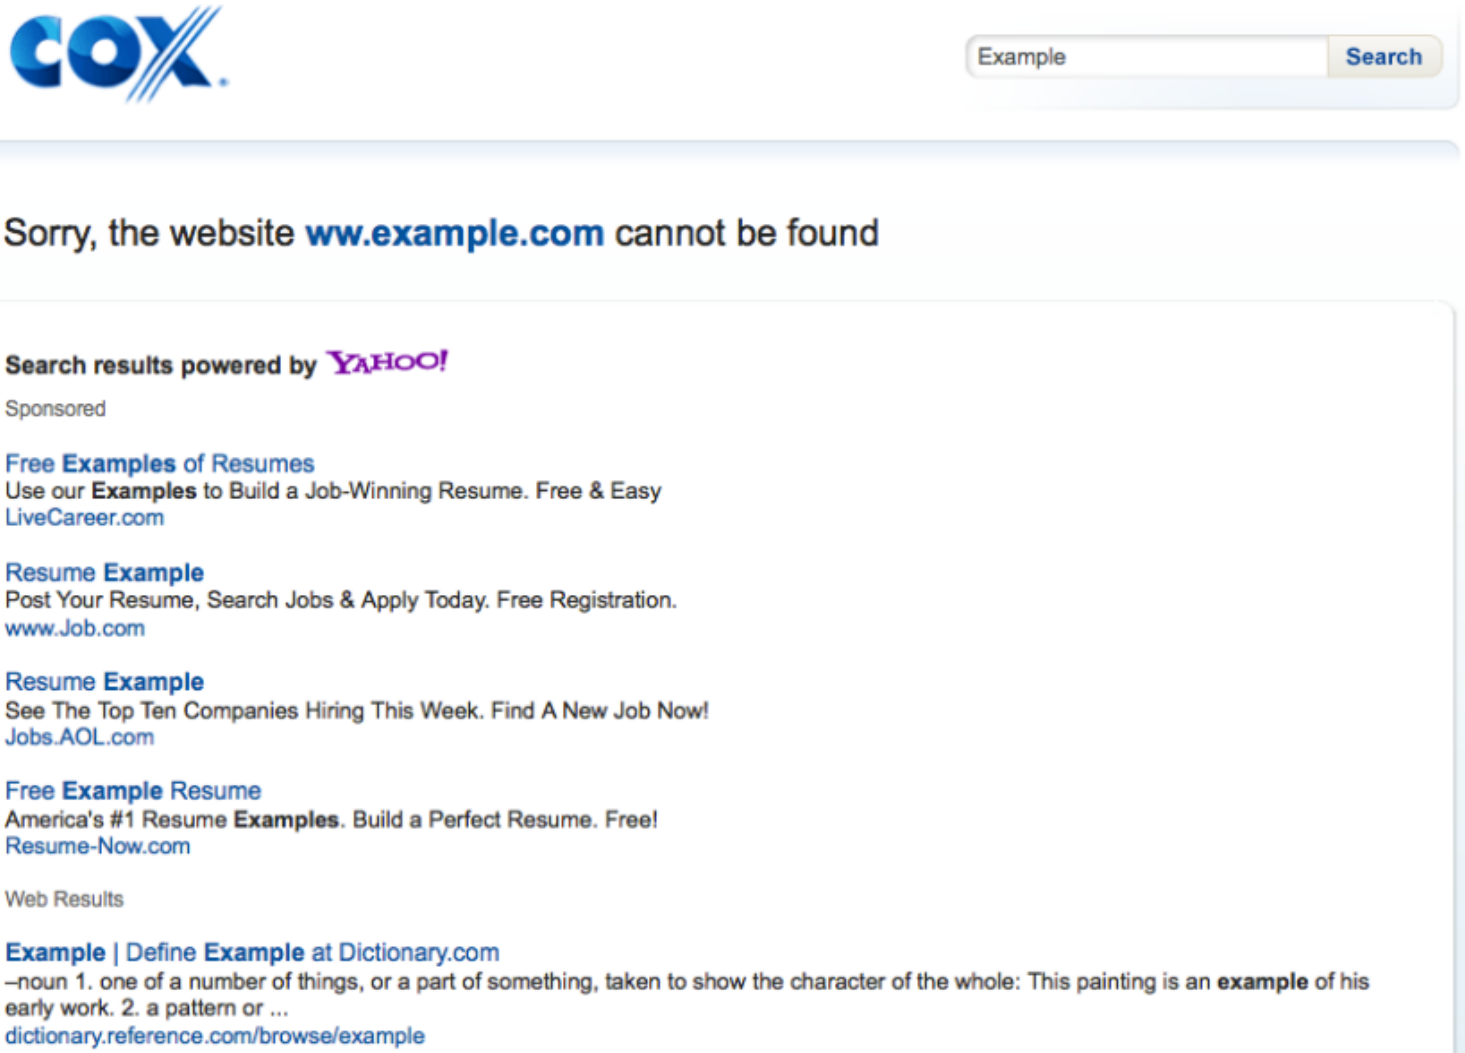

DDoS Attacks on Root Servers

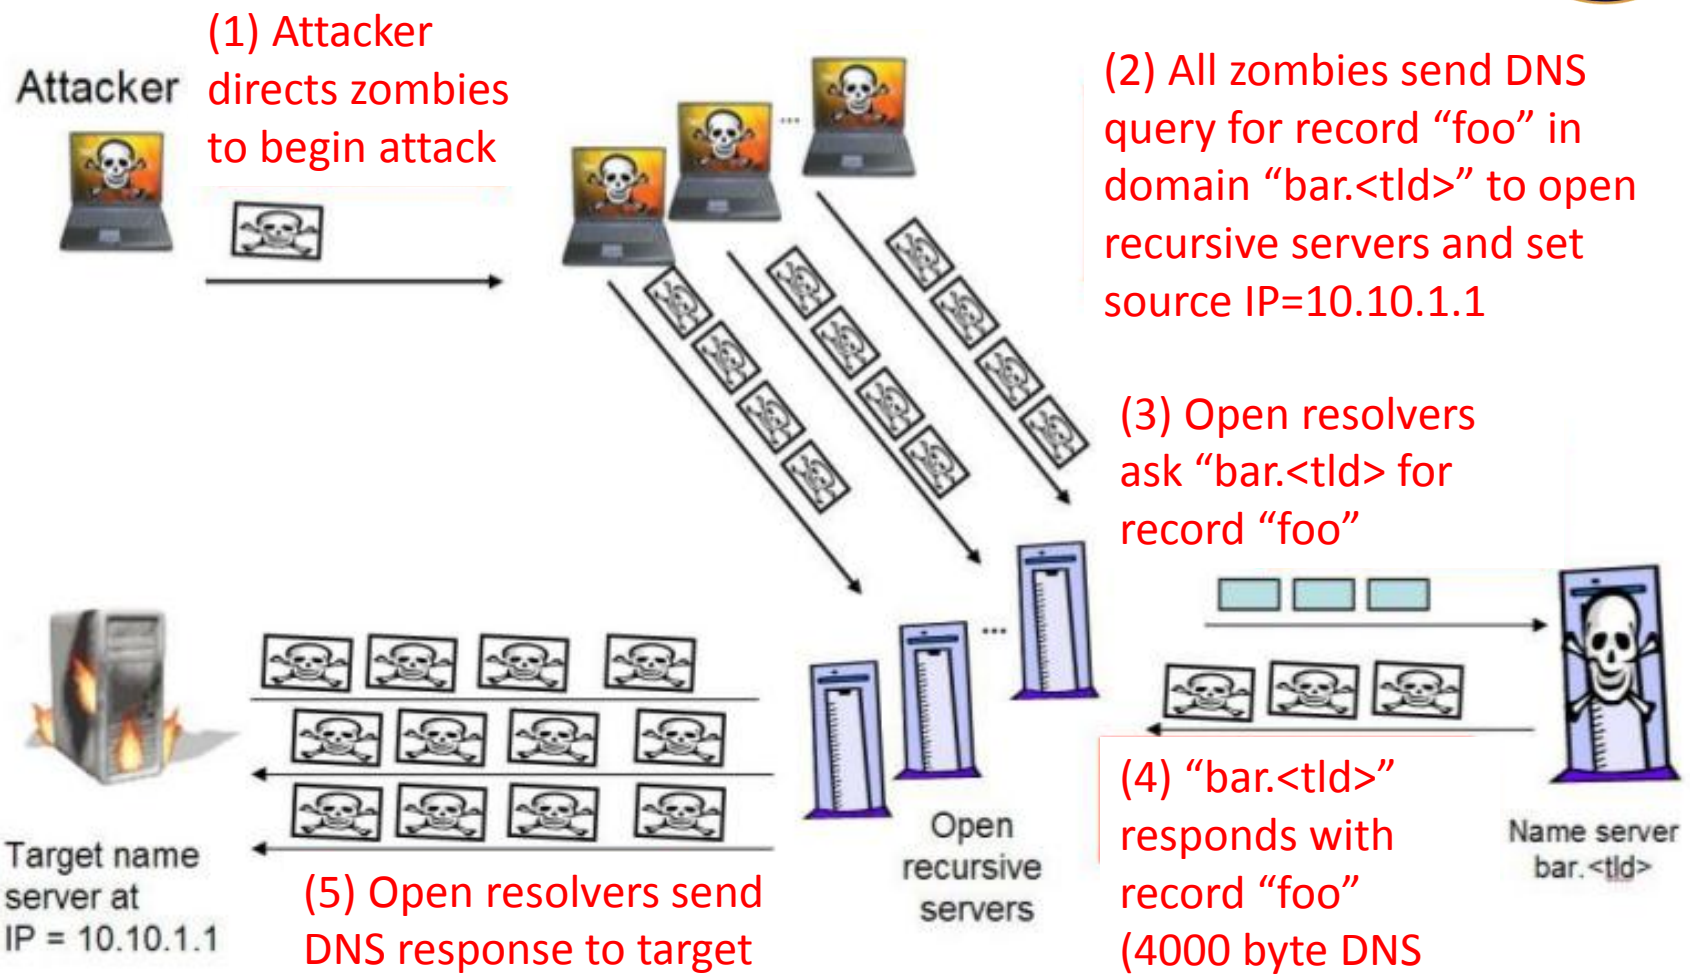

Other Problems

- Domain Squatting (cyber squatting)
- DNS Cache poisoning
- DNS ID Spoofing
- Mobility
- Dynamic DNS Updates
- Web DNS Updates
- Security